

REMARKS

Claims 5, 11 and 17 have been canceled by a prior amendment without prejudice or disclaimer of the subject matter thereof, while claims 2 and 7 - 18 have been canceled by the subject amendment without prejudice or disclaimer of the subject matter thereof. Applicants reserve the right to pursue the subject matter of any of the canceled claims in the subject application or subsequently filed continuing applications.

Claims 1 and 3 have been amended.

Claims 1, 3 - 4 and 6 are present in the subject application.

In the Office Action dated March 28, 2008, the Examiner has rejected claims 1 - 4, 6 - 10, 12 - 16 and 18 under 35 U.S.C. §103(a). Favorable reconsideration of the subject application is respectfully requested in view of the following remarks.

Initially, Applicants gratefully acknowledge the courtesies extended by Examiner Black during the telephone Interview of June 24, 2008. Applicants described one or more of the present invention embodiments and the manner in which the claims are distinguished from the cited art (similar to the manner discussed below). The Examiner indicated that further search and consideration would be required.

The Examiner has rejected claims 1 - 2, 4, 6 - 8, 10, 12 - 14, 16 and 18 under 35 U.S.C. §103(a) as being unpatentable over U.S. Patent No. 5,625,815 (Maier et al.) in view of U.S. Patent No. 5,659,701 (Amit et al.). This rejection is considered moot with respect to canceled claims 2, 7 - 8, 10, 12 - 14, 16 and 18.

Briefly, the present invention is directed toward a database management system installed in a data processing system and corresponding method. The database management system

manages a database having partitions for storing table data based on a partitioning schema, in which each partition has an associated partition identifier, and in which the database has database catalog information associated therewith. A partition identifier is identified in accordance with the partitioning schema, and the partition identifier is selected based on the contents of the query and the database catalog information. The query is executed against the identified partition. The technique improves the execution of queries while minimizing the consumption of network resources. For example, the technique eliminates the processing costs associated with directing the query to the appropriate partition, and transferring data from the data partition to a corresponding remote coordinator handling the query statement.

The Examiner takes the position that the Maier et al. patent discloses the claimed subject matter, except for an agent module. The Examiner further alleges that the Amit et al. patent discloses this feature, and that it would have been obvious to combine the Maier et al. and Amit et al. patents to attain the claimed invention.

This rejection is respectfully traversed. However, in order to expedite prosecution of the subject application, independent claim 1 has been amended and recites the features of: a database catalog accessible to each of the processing nodes and including information indicating data organization in the database, wherein the database catalog information includes partition maps associating table data with the networked partitions; the client query including a table name of the database table containing the desired data and a partition key value; a catalog cache to store within the network processing node information from the database catalog including partition map information for database tables; retrieving information within the database catalog and storing the retrieved information in the catalog cache, wherein the retrieved information

includes information identifying the table name and at least one networked partition containing the database table; analyzing the partition key value of the client query and partition map information retrieved from the catalog cache pertaining to the database table with the table name and containing the desired data to determine a specific partition; and executing the client query against the specific partition by determining the associated agent module for the specific partition based on the partition identifier and directing the determined agent module to execute the client query against, and retrieve the desired data satisfying the client query from, the specific partition.

The Maier et al. patent does not disclose, teach or suggest these features. Rather, the Maier et al. patent is directed toward a database computer system including a memory, residing in a plurality of interconnected computer nodes, for storing database tables. Each database table has a plurality of columns, a primary key index based on a specified subset of the columns, and an associated table schema. At least a subset of the database tables are partitioned into a plurality of partitions, each partition storing records having primary key values in a primary key range distinct from the other partitions. A transaction manager generates and stores an audit trail, each audit entry denoting a database table record event, such as an addition, deletion or alteration of a specified database table record in a specified one of said database tables (e.g., See Abstract).

Four online data definition procedures allow the structure of a database table to be altered while the database table remains available to execution of transactions, with minimal impact of the availability of the database table for transaction execution. The four online data definition procedures are a move partition procedure, a split partition procedure, a move partition boundary procedure, and a create new index procedure (e.g., See Abstract; Column 2, lines 39 - 46;

Column 7, lines 61 - 63; Column 9, lines 14 - 17; Column 10, lines 47 - 50; and Column 12, lines 19 - 22). Each of these online procedures has three or four phases of execution and begins with a user or operator entering a DDL statement specifying an alteration in the schema of a specified object, typically either a database table, an index or a partition (e.g., See Column 6, lines 9 - 13).

In a first phase, records of a table partition or the entire table are accessed using read only access, so as to generate a new partition, move records between two partitions, or to create a new index (e.g., See Abstract; Column 2, lines 47 - 57; Column 8, lines 15 - 29; Column 9, lines 41 - 50; Column 11, lines 10 - 20; and Column 12, lines 39 - 42). In a second phase, audit trail entries are used to clean up the data structures created during the first phase (e.g., See Abstract; Column 2, lines 58 - 62; Column 8, lines 30 - 50; Column 9, line 51 to Column 10, line 3; Column 11, lines 20 - 42; and Column 13, lines 12 - 22). In a third phase, access to the database table is briefly locked while audit trail entries created after the second phase are used to make final changes to the data structures created during the first phase, and while the database table schema is updated to reflect the changes to the database table produced (e.g., See Abstract; Column 2, lines 62 - 67; Column 8, line 51 to Column 9, line 10; Column 10, lines 4 - 40; Column 11, line 43 to Column 12, line 12; and Column 14, lines 11 - 30). In a fourth phase, used by the move partition boundary and split partition procedures, records in a database partition that are inconsistent with the modified database schema are deleted as a background operation while use of the database table by transactions resumes (e.g., See Column 3, lines 1 - 5).

Thus, the Maier et al. patent is directed toward modifying a database table's structure in accordance with a user command specifying the object (e.g., partition), while maintaining access

to database tables during the modification. There is no disclosure, teaching or suggestion of: networked processing nodes with agent modules each corresponding to a networked partition to establish a physical connection, execute a query and retrieve desired data from that networked partition; retrieving and storing information from a database catalog accessible to network processing nodes in a local catalog cache of a network processing node; determining the specific partition containing the desired data satisfying a query based on a partition key value of the query and partition map information retrieved from the catalog cache relating to the database table containing the desired data; and executing the query against that specific partition by determining the associated agent module for the specific partition based on the partition identifier of the specific partition and directing the determined agent module to execute the client query and retrieve the desired data from the specific partition as recited in the independent claim.

The Amit et al. patent does not compensate for the deficiencies of the Maier et al. patent. Rather, the Amit et al. patent is directed toward a multi-processor computer system executing a single-thread program having a plurality of callable routines. The multi-processor computer system makes it possible to convert a conventional single-thread program for operation on a multi-processor system without any significant modification to code. Although the Amit et al. patent refers to “partitions”, each of these partitions is associated with a corresponding processor, where procedures making up a program are allocated to the partitions (or processors) (e.g., See Column 6, lines 6 - 8). Thus, the partitions referred to within the Amit et al. patent are not database partitions as recited in the claims.

Amendment
U.S. Patent Application Serial No. 10/759,154

Since the Maier et al. and Amit et al. patents do not disclose, teach or suggest, either alone or in combination, the features recited in independent claim 1 as discussed above, this claim is considered to be in condition for allowance.

Claims 4 and 6 depend, either directly or indirectly, from independent claim 1 and, therefore, include all the limitations of their parent claim. The dependent claims are considered to be in condition for allowance for substantially the same reasons discussed above in relation to their parent claim and for further limitations recited in the dependent claims.

The Examiner has rejected claims 3, 9 and 15 under 35 U.S.C. §103(a) as being unpatentable over the Maier et al. and Amit et al. patents, and further in view of U.S. Patent Application Publication No. 2003/0233347 (Weinberg et al.). This rejection is moot with respect to canceled claims 9 and 15. Briefly, the present invention is directed toward a database management system installed in a data processing system as described above.

The Examiner takes the position that the Maier et al. and Amit et al. patents disclose the claimed subject matter, except for building a subset of the database catalog. The Examiner further alleges that the Weinberg et al. publication discloses this feature, and that it would have been obvious to combine the Maier et al. and Amit et al. patents with the Weinberg et al. publication to attain the claimed invention.

This rejection is respectfully traversed. Initially, claim 3 depends, either directly or indirectly, from independent claim 1 and, therefore, includes all the limitations of its parent claim. Claim 3 has been amended for consistency with its amended parent claim. As discussed above, the Maier et al. and Amit et al. patents do not disclose, teach or suggest the features of: networked processing nodes with agent modules each corresponding to a networked partition to

establish a physical connection, execute a query and retrieve desired data from that networked partition; retrieving and storing information from a database catalog accessible to network processing nodes in a local catalog cache of a network processing node; determining the specific partition containing the desired data satisfying a query based on a partition key value of the query and partition map information retrieved from the catalog cache relating to the database table containing the desired data; and executing the query against that specific partition by determining the associated agent module for the specific partition based on the partition identifier of the specific partition and directing the determined agent module to execute the client query and retrieve the desired data from the specific partition as recited in the claim.

The Weinberg et al. publication does not compensate for the deficiencies of the Maier et al. and Amit et al. patents. Rather, the Weinberg et al. publication is directed toward structuring, storing and retrieving data in database systems utilizing relational objects or qualifiers, and is merely utilized by the Examiner for an alleged teaching of a subset of catalog data. However, the catalog data relied on by the Examiner refers to items of manufacture or products (e.g., See Paragraph 0046), as opposed to database catalog information including partition maps associating table data with networked partitions as recited in the claim.

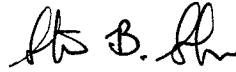
Since the Maier et al. and Amit et al. patents and Weinberg et al. publication do not disclose, teach or suggest, either alone or in combination, the features recited in claim 3 as discussed above, this claim is considered to be in condition for allowance.

Amendment

U.S. Patent Application Serial No. 10/759,154

The application, having been shown to overcome the issues raised in the Office Action, is considered to be in condition for allowance and a Notice of Allowance is earnestly solicited.

Respectfully submitted,



Stuart B. Shapiro
Registration No. 40,169

EDELL, SHAPIRO & FINNAN, LLC
1901 Research Blvd., Suite 400
Rockville, Maryland 20850-3164
(301) 424-3640
Delivered: 06/25/08